

Mentoring Undergraduates into Cyber-Facilitator Roles

Stephen Lien Harrell
Purdue University
sharrell@purdue.edu

Marisa Brazil
Purdue University
brazil@purdue.edu

Alex Younts
Purdue University
ayounts@purdue.edu

Daniel T. Dietz
Purdue University
ddietz@purdue.edu

Preston Smith
Purdue University
psmith@purdue.edu

Erik Gough
Purdue University
goughes@purdue.edu

Xiao Zhu
Purdue University
zhu472@purdue.edu

Gladys K. Andino
Purdue University
gandino@purdue.edu

ABSTRACT

There is a shortage of training programs for research cyber-facilitators and the need is only growing, especially in academia. This paper will discuss the importance of developing a workforce at the undergraduate level, creating a formal program for training and mentoring undergraduates in Research Computing at Purdue University, and how the approach to mentoring has evolved. The hands-on training and mentoring program has changed from one with students working as junior HPC administrators, performing hardware break-fix in a relative vacuum, to one with students working closely with their mentors, building real-world cyberinfrastructure solutions, such as distributed computing environments. More recently, the mentoring program has grown to include facilitating and supporting research applications with the Purdue user community. Finally, outcomes for the students in these programs lessons learned will be discussed.

CCS CONCEPTS

• **Social and professional topics** → **Informal education**; • **Software and its engineering** → **System administration**;

KEYWORDS

Student, Cyber-Facilitator, System Administration, HPC, Research Computing

ACM Reference Format:

Stephen Lien Harrell, Marisa Brazil, Alex Younts, Daniel T. Dietz, Preston Smith, Erik Gough, Xiao Zhu, and Gladys K. Andino. 2018. Mentoring Undergraduates into Cyber-Facilitator Roles. In *PEARC '18: Practice and Experience in Advanced Research Computing*, July 22–26, 2018, Pittsburgh, PA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3219104.3219138>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PEARC '18, July 22–26, 2018, Pittsburgh, PA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6446-1/18/07...\$15.00

<https://doi.org/10.1145/3219104.3219138>

1 INTRODUCTION

Comprehensive cyberinfrastructure (CI) has radically transformed the way scientific discoveries can take place. CI-powered methods and tools are applied to a wide range of disciplines to enable exciting research topics and create new opportunities for multidisciplinary collaborations. Efficient computing resources permit exploration of complex phenomena that can not be observed or replicated by experiment. Recently, data-intensive science has emerged as, the fourth paradigm of scientific discoveries [10]. Scientific instruments, simulations and sensor networks are generating more data than ever before, and the rate of increase is accelerating rapidly. This visualization technology contributes to the understanding of data across disciplines by opening up insight into the complex relationships that exist within the data.

Many campuses have invested heavily in advanced cyberinfrastructure, which encompasses high-performance computing and storage systems and fast networks, among other resources for their research enterprises. A critical aspect of these efforts is to provide expert support for efficient and effective workflows on these state-of-art systems, from cyber-facilitators to domain researchers. However, effective cyber-facilitators require training on a wide range of topics [12], which are not normally covered by an undergraduate curriculum [9]. This is confounded by the fact that cyber-facilitation skills often encompass other fields, which can cause regular lateral movement out of the research computing field without a corresponding movement back in. Additionally, recruitment of undergraduate students targets, but is not limited to, groups underrepresented in the cyber-facilitator field, including minorities and women.

2 STUDENT EMPLOYMENT AS CYBER-FACILITATORS

At Purdue it has been found that student employment is one of the most effective ways to mentor students. It engages students professionally and gives them a reason to learn. Additionally, providing a path to more interesting work, along with the promise of an upgraded position, can also be a powerful motivator for them to learn skills that are crucial in the organization.

2.1 Hardware Break-Fix

Purdue purchases a five to six hundred node cluster every year and runs them in production for five years, minimum. For almost a decade, we have run at a stable count of around 3,000 compute nodes, with a cluster being decommissioned and a new one being built every year. It is a lot of hardware to simply shuffle, let alone keep under repair.

Working under the direction of full-time systems administrators, most students begin their careers performing hardware maintenance. They learn the ins and outs of replacing memory, hard disks, and power supplies while working with our hardware vendors. To keep track of work and provide data for failure analysis, a custom built ticketing and issue-tracking system is employed. Students have their work dispatched automatically from monitoring systems and provide details of fixes that are recorded for each node. The daily break-fix provides a gentle but structured introduction to HPC systems administration. The procedures tend to be pre-canned and, as changes are needed, the focus is on good documentation skills.

Once a student has done their time in the hardware trenches, they move to specialized deployment projects. These can include special purpose clusters, life cycle replacing back-end servers, and similar projects which include a heavy hardware component. These projects also introduce the students to the configuration management system and our change management processes.

2.2 System Administration

Once a student has acquired skills resembling a junior HPC systems administrator, they find themselves being given public, customer-facing projects to operate. The Purdue Recycled Clusters was one such long standing, student run project. As the computer labs around campus were replaced, the research computing division acquired the castaway desktops. This hardware pool was incorporated by students as the very first Linux cluster on campus and was primarily run by the student junior administrators. Over the course of nearly two decades, the Recycled Clusters provided millions of free compute hours to the campus and reached almost a thousand nodes at its peak. Beyond the Recycled Clusters, students are usually involved in up and coming technology deployments. The very first Purdue scientific cloud and later Teragrid resource, Wispy, was led by a student. As well, the very first Hadoop clusters on campus came out of student projects.

While full time staff focus on keeping flagship resources available and providing immediate value on projects, student junior systems administrators are free to explore riskier areas. Not all exploratory projects get off the ground, but when looked after properly, they provide a safe area for students to cut their teeth and learn important lessons that translate into hireable talent.

The results are trained students, who are well rounded in cluster hardware, software management, and Purdue Research Computing's processes.

2.3 Software Development

Undergraduate students have been utilized for various software development projects. Outlined in this section are several examples

of development projects our undergraduate students have participated in, and how these projects benefited both the student and the institution.

2.3.1 User Support Tools. To assist the users of our HPC resources, we have created several tools which are installed on compute resources. Undergraduates have spearheaded the development of several of these software packages, providing invaluable capability and directly exposing them to HPC resources.

One example is a tool that provides a summary of job submission queues, and jobs submitted and running in those queues. This tool was primarily developed by an undergraduate student under the guidance of a full-time staff member. The project exposed the student directly to PBS and job management systems through programmatic manipulation of PBS commands. Such exposure allowed the student to get a taste of some of the intricacies of such systems. For example, the student had to be familiar with node and job scheduling policies, neither of which are simple concepts.

2.3.2 Web Portal. A core component to our operations is a web-based portal that allows faculty and researchers to manage access to their HPC resources [3]. This project began as a staff-led project, but over the years has had several undergraduate students work on enhancements, bug fixes, and new features to the web portal. This work exposes students to all facets of our HPC operations, as this portal touches nearly every aspect from user management, queue management, storage management [4] to an integrated, custom-built customer relation management tool.

2.3.3 Kickstand. Like the previously mentioned web portal, another web-based tool was developed to manage compute node inventory and began as a staff-led project. Undergraduate students worked on this project to add functionality. Recently, another set of undergraduate students have been working on a project to rebuild this system using their experience with the original system to build an improved version. Work on this project has allowed the undergraduates to gain exposure to what goes into managing a compute cluster. This tool tracks every compute node in regards to its hardware, repairs, and the software configuration on the compute nodes so the students are exposed to many facets of HPC operations. While this project is largely web programming, the integration with HPC system tools is especially relevant in mentoring students in system administration and facilitation.

2.4 Cybersecurity of Research Networks

In a project sponsored by the Cybersecurity Innovation for Cyberinfrastructure (CICI) program of NSF, Purdue is integrating state-of-the-art network monitoring and intrusion detection tools within the entire campus research cyberinfrastructure, which is designed for open science, to meet growing demands from the increase in industrial, governmental and global partnerships. This project provides undergraduate/high school students a unique opportunity to participate in real-world implementation in deployment and operation of networking hardware and software. During the two-year project life, the students serve as project interns and are fully integrated into the project from design through implementation and evaluation. The team consists of five men and one woman, one of which is a high school junior, who was recruited from the 2017

Student Cluster Competition team of William Henry Harrison High School.

The wide technical scope of this project requires that the group of students as a whole be knowledgeable in networking, network security and systems administration. During recruitment we found that students' experience (and excitement) was usually centered around one of these areas. Students were chosen based on their expertise and strengths across all areas, while ensuring at least one strong candidate was chosen in each specific area. It is our hope that throughout the life of the project, the students who strongly associate themselves with one area (i.e. networking) will find they have gained experience and new knowledge from the other areas as well.

Following the Kolb's cycle of experiential learning [11], each student intern is paired with a mentor from the project team to help them hone skills, answer questions, and guide student reflection and synthesis of new knowledge. It is not the mentor's role to make decisions on the technical aspects of the project or to provide a simple project plan for the students to implement. Instead, the mentor acts as a guiding hand to lead the students in properly and fully assessing the technical problems of the project and to help them arrive at appropriate solutions. Decision making is left in the hands of the students, and the sharing and discussion of alternate opinions, methods and ideas is encouraged.

To facilitate this discussion, the project and its implementation are outlined in a Confluence wiki. The wiki acts as a central knowledge base, allowing each student to continually document their work as they go from understanding a problem and performing research, to deciding and implementing a solution. Additionally, the technical details of the project are discussed in a bi-weekly open forum attended by all interns, mentors and project staff. Students are required to prepare slides to present the technical problem they are addressing, the research they have done on the topic, and implementation details. By consistently having students present their work, we hope to eliminate any black boxes in the final implementation and ensure students are exposed to all parts of the project, even those they are not directly accountable for.

The mentor approach and open forum discussions have produced a desirable effect on student/mentor interaction. During initial forums, students directed questions to the mentors looking for immediate answers. Questions like "How are we going to implement X?" or "What is our strategy for configuring Y?" were common. Once the students realized that the mentors were not offering a defined project blueprint, their behavior changed. Students are now using critical thinking skills and synthesizing knowledge from different sources to come up with answers to technical problems and, more importantly, articulating the reasons that led them to their conclusion.

A rigorous student assessment plan was developed to determine levels of student motivation, self-directed learning, engagement, and career interest. Assessment of the students are conducted in order to:

- **Assist Learning:** students are regularly assessed through their guided reflections with their mentors and provided with immediate feedback to help them build their knowledge and skills

- **Measure Individual Achievement:** mentors use a rubric to provide short monthly reports on how well students are engaging in practice, building skills, and applying knowledge
- **Foster Professional Development:** Communication, leadership skills, time management, and teamwork are all goals

The assessor on the project uses interviews conducted with students every three months to determine levels of student motivation, self-directed learning, engagement, and career interest.

It is our hope that by exposing the students to multiple facets of an advanced cyberinfrastructure monitoring project, while giving them the freedom to drive the implementation with their own ideas, they will gain confidence in their decision making skills and will acquire valuable experience in the research and design of advanced computing infrastructure. While the project is still in its early stages, we feel that the mentors and project staff have laid solid groundwork for both student growth and project success.

2.5 Facilitator Internships

Facilitator internships are an important part of our mentoring strategy. The goal of these internships is to teach the skills necessary to be a full stack research cyber-facilitator. This goal has two main thrusts, HPC Systems Administration and Scientific Application Support. A candidate that completes this internship will be prepared to enter an organization that has no HPC, build an HPC cluster, support that cluster and its users who run scientific codes. In order to enable this professional goal, we split the internship into two time periods where the interns are embedded in our HPC System Administration and Scientific Application Support groups. In those groups, they learn a variety of skills listed below.

2.5.1 System Administration for Research. After the intern is embedded within the research administration team the candidate will be able to:

- Create a basic design for a typical HPC cluster and Big Data cluster
- Create a hardware parts list for that cluster design
- Assemble the hardware into a working cluster
- Install a base OS onto all nodes in the cluster
- Compile and install a list of software needed for the operation of an HPC cluster such as a scheduler, common scientific libraries, etc. or a Big Data cluster such as Hadoop, Spark, MapReduce, etc.
- Install and maintain services for monitoring cluster such as Ganglia, Sensu, Nagios, etc
- Configure the installed machine to work as a cluster making sure ssh keys are in place, proper scheduling rules are in place, etc.
- Use a configuration management system such as xCAT, CFEngine [2] or Puppet to configure and maintain the cluster
- Create tests and monitors to confirm the cluster is in good health from a node and service perspective
- Take the entire cluster down for maintenance, do the maintenance, and return the cluster to service within allotted time

2.5.2 Scientific Application Support. After completing the faculty-facing portion of the internship, the candidate will be able to:

- Build a basic tool-chain including the MPI libraries for researchers to base their science codes upon
- Compile and package complex scientific codes with many dependencies
- Create and use environment modules using LMOD [13]
- Answer users' questions about the use of research clusters
- Answer users' questions about the use of scientific software
- Talk with users about their problems and communicate solutions in an understandable and relatable way
- Communicate planned and unplanned service disruptions
- Assist with performance benchmarking
- Assist with the training of faculty and graduate students to use research clusters

2.6 External Internships as Funding Sources

Various departments at Purdue that focus on student success provide avenues for funding internships. These are great opportunities to find extra funding and enhance the existing student programs run by the research computing division.

2.6.1 Summer Stay Scholars. The Purdue Summer Stay Scholars program provides tuition reimbursement for students to stay on campus over summer break. These students take three credit hours of courses and work 140 hours for a campus unit. The goal is to entice departments to trial potential students and it is seen as a success when a student is still working even after their initial summer experience.

During the first trial of this program, one student worked to build machine learning and artificial intelligence containers for Singularity and the supercomputers on campus. Successive students have gone on to build automation frameworks and perform improvements that languished due to a lack of attention from full time staff.

2.6.2 Executive Vice President and Treasurer Internship Program. For two summers, the Treasurer of Purdue funded departments reporting to that office to employ full-time summer interns. The goal was to provide an opportunity for departments to find and recruit graduating talent. The program was medium-sized at forty students, five of which were brought into the Information Technology group.

The single student who interned for the Research Computing division worked to benchmark current generation HPC interconnects for the Coates supercomputer. This student was given access to a small cluster of 8 nodes and a stack of various options: SDR and DDR Infiniband, plus 10Gbps iWARP. The results of this program became the cornerstone of several clusters that came afterwards.

Upon graduation, the student intern accepted a position in Research Computing and has stayed in HPC.

3 NON-EMPLOYMENT CYBER-FACILITATION MENTORSHIP

Research Computing at Purdue is involved in a number of activities with students beyond student employment. They support student learning and give more opportunities for staff to mentor undergraduates.



Figure 1: Song-Hui Ryu, student recipient of a WHPC travel grant to PEARC17, presenting about her experience at August 30, 2017 WHPC meeting.

3.1 Women In HPC

Purdue Research Computing founded our Women in HPC (WHPC) program in the fall of 2016, as part of an IT broader engagement initiative intended to expose undergraduate and graduate students to strong female role models and encourage women to pursue research and careers in high performance computing (HPC) and technology-related fields. The WHPC membership is comprised of undergraduates, graduate students, staff and faculty women and men at Purdue University.

Why does Purdue value developing a diverse workforce? According to the Bureau of Labor and Statistics, there were about "8.6 million STEM related jobs in May 2015, representing 6.2 percent of U.S. employment" and "computer occupations made up nearly 45 percent of STEM employment" [5]. However, according to the National Center for Women in Technology (NCWIT), the number of women in computing occupations has decreased 10 percent since 1991, with only 26 percent of women in these fields [1]. Although that has slightly improved since then, the numbers are still dismal.

Research indicates a variety of benefits to increasing diversity in the workplace, including: improved collaboration, innovation and productivity in teams due to the exchange of new ideas and perspectives, as well as increased financial performance for companies. In fact, some studies have shown that diverse teams consistently perform at a higher level [1]. In our effort to address the lack of women in HPC and other technical fields, Purdue WHPC program's mission is to promote awareness of HPC within our local Purdue University campus by providing a supportive community that provides a sense of membership, inclusion and belonging, and a space for female students to develop their interests in this field. The WHPC program offers students a variety of activities and resources including:

- presentations from invited speakers, technical experts and role models (about 5-6 per year);
- promoting participation at technical conferences through travel grant scholarships;
- exposure to HPC-related topics, research and careers;

- professional and personal development opportunities (e.g. public speaking and presentations of posters and papers);
- development of leadership, networking, communication, and social skills; and,
- the opportunity to build a community with students, faculty and staff with similar interests.

The WHPC staff is also in the process of implementing a mentor program. The goals of the mentor program are to:

- Create additional professional development and networking opportunities with experts and role models in the field;
- Keep skills current; and
- Exchange ideas and discover different research fields or areas of interest.

Purdue University recognizes the importance of diversity and has taken a lead in promoting the advancement of women in HPC and other historically male-dominated fields. By educating a diverse workforce on the benefits of advanced computing resources, we are developing future facilitators and researchers with access to, and knowledge of, HPC resources who can, in turn, share this valuable knowledge with their community of peers.

3.2 Student Cluster Competition

Mentoring undergraduates in HPC is core to one of the most interesting events at the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC) every year, the Student Cluster Competition (SCC) [7]. Students build and run a small HPC cluster and then use that cluster to run a variety of scientific and engineering applications. Each cluster is constrained by 3000 watts and the competition runs for approximately 48 hours straight at the SC conference, however, other SCC events have slight variations on the time frame.

This type of competition provides a breadth-first introduction to facilitation and many of the hot-topics in the field such as HPC systems administration, scheduling, application building and usage on parallel platforms. Students' understanding of how many modern scientists operate from a full-stack perspective is key. Interconnect and storage problems, operating system and scientific application problems and data management problems are all critical to succeeding in the competition and most of the mentorship to help students start to understand these pieces starts at the team's institution.

At Purdue, the staff from Research Computing have mentored 14 teams (including one high school team) in various competitions and over 60 undergraduate students since 2007 [8]. Research Computing staff have used a variety of methods for inspiring student participation, including for-course credit, after hours extra curricular activities and direct one-on-one mentorship. However, most of these methods come down to a core set of guidelines: Create an interconnected problem for the students to solve, give them time to explore the problem, and then guide them when they run into trouble. Given the right problem set, this is a very effective way to teach what matters for the competition and facilitation.

Purdue has also pioneered partnering with new institutions to provide access to HPC mentoring to students at institutions that are underserved in the HPC ecosystem. This has been a success both for Purdue and the partner institutions. Working on teams



Figure 2: Purdue Team at the Student Cluster Competition at SC10.

with students from different background and different goals can bring new understanding of the impact and draw of HPC.

The SCC program has been a success at Purdue in two ways. We have been able to use this program to gauge interest in HPC and recruit students who are interested in learning more about HPC, post-competition. Additionally, the SCC is a good tool to extend current students' knowledge of HPC, enabling us to move them from hardware break fix to junior administrator and facilitation roles. Additionally, many students who have attended the SCC have also gone on to cyber-facilitation roles at Purdue and elsewhere. For instance, in Fig 2, two undergraduate students on this team are current Research Computing staff at Purdue.

3.3 In the Classroom

In concert with interested faculty, staff in research computing have taught facilitation skills in domain science courses around the university. These collaborations can range from basic scientific application use to full-stack courses ranging from one day to the entire semester.

The undergraduate Scientific Computing course in the department of Earth, Atmospheric and Planetary Sciences is one example of a successful course that focused heavily on facilitation at Purdue. In this class, students were each given four desktop machines that were recycled from the on-campus computer labs. Students installed, networked and configured these machines to run the Weather and Research Forecasting model (WRF) [14]. The following topics were covered over the semester:

- History of Numerical Weather Prediction
- Introduction to High Performance Computing
- Hardware setup and configuration
- Installation automated and manual installation of Linux
- Schedulers
- Interconnects and Storage Hierarchy
- Shared File Systems
- Cluster configuration
- Application (WRF) and library installation
- Configuration and running of WRF using MPI
- Introduction to Python with an emphasis on WRF output

By the end of the class, students were able to run WRF across their four machines and generate an image of the forecast, as well as pull out specific data from the output of the model. This is the most extreme example of cyber-facilitation education in the classroom administered by Research Computing. These basic skills are critical for not only the supporters of science but for budding scientists themselves, which makes teaching a subset of these topics very attractive for faculty who teach the use of numerical models.

4 REACHING BEYOND THE CAMPUS

Since 2013 Purdue Research Computing staff has been mentoring students from Universidad EAFIT. EAFIT has created one of the first HPC centers in Colombia. However, cyber-facilitators are rare in the country, so EAFIT has committed to training their students in cyber-facilitation by necessity. In order to help EAFIT meet their need, Purdue Research Computing staff has created an internship for top students in the Apolo center at EAFIT. This program fits well within EAFIT, since the school requires each undergraduate complete a six month internship. Students come to West Lafayette and participate in our Facilitator Internship, spending half their time working closely with the research system administration team and then research support. Additionally, staff at Research Computing at Purdue have helped them in this effort by providing training on-site [6]. Since then?, Purdue joined teams with EAFIT, to form a combined SCC team for SC14 and ISC15. In 2015, EAFIT began bringing its own team.

5 LESSONS LEARNED

During the course of our student mentoring program, we have learned that physical proximity is a key factor in the success of mentoring a student. The most effective outcomes occur when students are co-located with full-time staff, and work closely on a daily basis. When students are geographically isolated from their mentors, the relationships are difficult to maintain, and learning opportunities are limited. Several methods of supervision have been attempted, from a dedicated "student coordinator" to mentor/mentee pairings, with the mentor/mentee relationship producing better-trained students.

Another indicator of engagement in facilitation after student employment, unsurprisingly, is the projects and extra curricular activities students were involved in. Although Purdue Research Computing has need for both software development and research system administration, students who are engaged in software development tend not to go on to pursue facilitation; on the other hand, there has been good retention with our junior administrator and student cluster competition initiatives.

While software developers and system administrators have been successful in developing future facilitators, data points on students hired into a purely user-facing role are limited. For this reason, the internship program described in 2.5.2 specifically includes a work rotation in application support.

We have also identified opportunities for improving our value to the students themselves. For example, while Purdue may not always have an open position in which to hire a student, our professional network within the national community can provide help placing students in facilitator roles elsewhere. Providing placement

assistance for students interested in work in research computing would be a valuable addition to this program.

6 REAL WORLD SUCCESS STORIES

Over the course of 14 years, approximately 70 students have passed through the Purdue Research Computing student program, with nine students converted into full time staff at Purdue. Additionally, four students have gone on to work professionally in the research computing community nationwide – a conversion rate of 20% entering the research computing community.

Research computing students have become HPC system administrators, research programmers or departmental facilitators at Purdue. Other students have performed similar roles at other major research universities, national labs, and NSF computing centers. Additionally, they have gone to work at startups in cloud-based high performance computing, as well as HPC hardware and storage vendors.

7 CONCLUSIONS

Mentoring students with hands-on, experiential learning as a tool for creating cyber-facilitators has proven to be a successful technique to develop a research computing workforce at Purdue University. Additionally, extra curricular programs such as the Women in HPC program and student cluster competitions complement and replenish the ecosystem of undergraduates engaged in HPC. The program at Purdue of advancing student capabilities along a trajectory from hardware technicians to developers and junior system administrators, (with students coming from work study, internship programs, or NSF projects) has produced 14 research computing professionals since 2004. Augmenting the professional staff of an academic HPC center with student admins is a successful and cost-effective way to meet two university missions of education and discovery at the same time.

8 ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1738981.

REFERENCES

- [1] Catherine Ashcraft, Brad McLain, and Elizabeth Eger. 2016. WOMEN IN TECH: THE FACTS. (2016).
- [2] Mark Burgess et al. 1995. Cfengine: a site configuration engine. In *in USENIX Computing systems*, Vol. Citeseer.
- [3] Kevin D. Colby, Daniel T. Dietz, Preston M. Smith, and Donna D. Cumberland. 2014. Self-service Queue and User Management in Shared Clusters. In *Proceedings of the First International Workshop on HPC User Support Tools (HUST '14)*. IEEE Press, Piscataway, NJ, USA, 22–31. <https://doi.org/10.1109/HUST.2014.5>
- [4] Daniel T. Dietz, Lev A. Gorenstein, Gregory S. Veldman, and Kevin D. Colby. 2017. Shared Research Group Storage Solution with Integrated Access Management. In *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact (PEARC17)*. ACM, New York, NY, USA, Article 14, 7 pages. <https://doi.org/10.1145/3093338.3093354>
- [5] Stella Fayer, Allan Lacey, and Audrey Watson. 2017. STEM Occupations: Past, Present, And Future. (January 2017).
- [6] Stephen Harrell. 2015. Construcción de clústers para computación científica. Presented at Escuela de Verano Universidad EAFIT.
- [7] Stephen Lien Harrell, Hai Ah Nam, Verónica G. Vergara Larrea, Kurt Keville, and Dan Kamalic. 2015. Student Cluster Competition: A Multi-disciplinary Undergraduate HPC Educational Tool. In *Proceedings of the Workshop on Education for High-Performance Computing (EduHPC '15)*. ACM, New York, NY, USA, Article 4, 8 pages. <https://doi.org/10.1145/2831425.2831428>

- [8] Stephen L. Harrell, Preston M. Smith, Doug Smith, Torsten Hoefler, Anna A. Labutina, and Trinity Overmyer. 2011. Methods of Creating Student Cluster Competition Teams. In *Proceedings of the 2011 TeraGrid Conference: Extreme Digital Discovery (TG '11)*. ACM, New York, NY, USA, Article 50, 6 pages. <https://doi.org/10.1145/2016741.2016795>
- [9] Peter J. Hawrylak, John Hale, and Mauricio Papa. 2017. Undergraduate Educational Pathways for Developing a High-Performance Computing Workforce. In *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact (PEARC17)*. ACM, New York, NY, USA, Article 53, 4 pages. <https://doi.org/10.1145/3093338.3093355>
- [10] Tony Hey, Stewart Tansley, Kristin M Tolle, et al. 2009. *The fourth paradigm: data-intensive scientific discovery*. Vol. 1. Microsoft research Redmond, WA.
- [11] David A. Kolb. 1984. *Experiential learning: Experience as the source of learning and development*. Prentice-Hall, Englewood Cliffs, NJ.
- [12] Scott Lathrop. 2016. A Call to Action to Prepare the High-Performance Computing Workforce. *Computing in Science and Engg.* 18, 6 (Nov. 2016), 80–83. <https://doi.org/10.1109/MCSE.2016.101>
- [13] Robert McLay, Karl W. Schulz, William L. Barth, and Tommy Minyard. 2011. Best Practices for the Deployment and Management of Production HPC Clusters. In *State of the Practice Reports (SC '11)*. ACM, New York, NY, USA, Article 9, 11 pages. <https://doi.org/10.1145/2063348.2063360>
- [14] William C Skamarock, Joseph B Klemp, Jimy Dudhia, David O Gill, Dale M Barker, Wei Wang, and Jordan G Powers. 2005. *A description of the advanced research WRF version 2*. Technical Report. National Center For Atmospheric Research Boulder Co Mesoscale and Microscale Meteorology Div.